

METHOD AND APPARATUS FOR NOISE FLOOR ESTIMATION

Field of the Invention

The invention relates to the field of telecommunications and in particular the signal processing in receivers of communication signals.

5 Background

In the art of telecommunications, it is desirable to estimate the average noise power contained in a received signal. Typically, the received signal is a frequency division multiplexed composite signal containing several distinct signals. It is important to be able to estimate the noise power without a priori knowledge of the number, nature, carrier-frequency, 10 bandwidth or amplitudes of the distinct signals.

In the art, noise is considered an additive phenomenon. It distorts a signal by adding random time-varying values to the amplitude of the non-noise component of the signal. It is thus of great importance to develop a radio receiver that is robust to the noise phenomenon.

The general problem of estimating the variance of the additive noise in a communication 15 signal composed of many narrowband frequency division multiplex signals is important in many radio communication systems. It allows setting the detectors threshold to its maximum sensitivity while limiting the probability of a false alarm to a value less than or equal to a desired level. The detector decides that a signal is present when the energy of a signal, over a period of time, exceeds a threshold value. The detected signals can then be passed to a 20 receiver or to measurement equipment, for demodulation or to acquire more information about the signal.

Currently there are several techniques that are used in this pursuit, but all come with sacrifices. One such technique is to reserve a portion of the bandwidth of the composite

signal to carry a "noise channel". Thus by measuring the power in this reserved band, an estimate of the noise power in the other channels can be obtained. It is common though for these "noise channels" to be difficult to identify as their placement in the spectrum of the composite signal is not governed by any global convention which leads to different channel assignments based on geographical location of the measurement equipment. Other errors are introduced when the "noise channel" is affected by the signals in adjacent channels. This results in a biased estimation of the noise floor. If the channel assignment is unknown, it may not be possible to identify the "noise channel" and determine the noise floor power level. In addition to these problems the noise floor tends to fluctuate with environmental and operation conditions.

In the art, it is common practice to set a noise threshold value, determined in large part by the estimated noise floor, whereby a signal is considered present if and only if it is above the noise threshold.

The traditional approach is to set a threshold value depending on the user needs. When high sensitivity is required, the threshold is set to a value close to the user observed noise floor level. When only strong signals are of interest, the threshold is set to a high value relative to the observed noise floor level. This approach is highly subjective and may result in being too aggressive or too conservative for the actually measured signal environment. The threshold is typically set manually in a communications system. This is not feasible when fast response to dynamically changing conditions is required. Very few automated techniques have been studied to address the general problem of rapidly and dynamically estimating the noise floor level or the signal-to-noise ratio. One typical method is to isolate a channel and to assign this channel as a noise only channel for estimation purposes. This method is difficult to apply in many new digital communication systems where the resources are often utilized close to their capacity. It is also a problem for dynamic signal environments based on multiple access. Another typical method is to assume that the composite wideband signal non-noise subspace dimension is known. This is again difficult to realize in practical situations with unknown signals where robustness is a key feature. On another front, a technique based on morphological binary image processing operators (similar to rank-order

filtering) as described in "Automatic Noise Floor Spectrum Estimation in the Presence of Signals", M.J. Ready, M.L. Downey, and L.J. Corbalis, *Proceedings of the 31st Asilomar Conference on Signals, Systems and Computers*, vol. 1, pp 877-881, November 1997, processes a binary image of the wideband received power spectrum of a composite signal has
5 been proposed. Thus, the approach does not process the time samples or frequency domain transformed data directly, but instead it processes an image of the spectrum. This involves two dimensional signal processing that is usually complex. Furthermore, practical performance results have yet to be quantified with such an approach. In general, techniques based on textbook results tend to assume several conditions that are not truly present in many
10 bands of the radio spectrum.

The computational complexity of a method is a limiting factor of the method's applicability. Typically, a spectrum monitoring receiver cannot observe the full spectrum it is allotted, thus it divides the spectrum into smaller blocks that can be scanned, each of these blocks possibly contains a plurality of signals. Each block is received sequentially, an estimate of the noise floor is arrived at, and then the receiver shifts its frequency to receive another block. This cycling through the allotted bandwidth must be done rapidly to ensure that all messages being transmitted in the different parts of the allotted spectrum can be received. If the method to derive a noise floor estimate is too complex then it either cannot keep pace with the shifting of the block or it causes the shifting of the blocks to slow down,
15 20 in the event of either result the method ceases to be useful.

Thus, it would be desirable to design a system that resolves the problems that the above mentioned methods have with narrowband signals of unknown nature, as described in "Noise Floor Estimation for Wide-Band FFT Filter Banks". It would also be desirable for a new system to have fully digital characteristics so as to permit real-time signal performance,
25 without requiring the dedication of a reserved frequency band for use as a noise channel.

Summary of the Invention

It is an object of the invention to remedy numerous problems found in the prior art
5 methods.

It is another object of the invention to provide a noise floor estimation method and apparatus capable of operation on a wideband signal without full knowledge of the nature of the narrowband signals contained therein.

10 It is another object of the invention to provide digital characteristics to allow real-time signal processing with low levels of computational complexity.

The preferred embodiment of the method of the invention, with very little a priori knowledge of the signals, allows both robust performance in non-ideal environments, and an automated maximum sensitivity threshold setting method. These modifications make the process of the invention an attractive solution to problems emerging in the field of digital
15 communications. As an added benefit, information derived by the method of the invention may be used in spectrum quality indication, interference analysis, and signal demodulation.

In the preferred embodiment of the method in accordance with the present invention, the noise floor is estimated by carrying out the steps of:

- a. representing the continuous signal as a series of discrete frequency and amplitude values;
- 20 b. creating a histogram based on the discrete frequency and amplitude values;

- c. deriving a noise floor estimate from characteristics of the histogram.

The discrete frequency and amplitude values are preferably obtained by the steps of:

- d. sampling a received wideband signal by a plurality of analogue-to-digital converters to generate a series of output signals;
- 5 e. windowing the output signals of the analogue-to-digital converters;
- f. applying a mathematical transform, preferably a Fast Fourier Transform, to the results of the windowing to obtain a series of discrete frequency values;
- 10 g. converting an amplitude of each discrete frequency value to log-domain representation; and
- h. rounding the log-domain representation of the amplitude for each discrete frequency value to the nearest integer value to generate a discrete amplitude value.

The windowing step preferably includes the steps of:

- i. selecting a discrete weighting function;
- 15 j. multiplying the value of each output signal of the series by a corresponding element of the discrete weighting function.

The amplitude of each value of the discrete frequency series is preferably converted to log domain representation by multiplying 20 by the base 10 logarithm of the magnitude of the element; whereby the log domain representation of the amplitudes results in the amplitudes being expressed either as decibel (dB) values, or as decibel milliwatt (dBm) values.

In a first preferred embodiment of the steps of creating and analyzing the histogram to estimate the noise floor, herein referred to as the discrete level crossing rate method, a histogram chart is constructed where the minimum value is equal to the minimum integer value in dB of the received signal strength, and the maximum value is equal to one less than the highest integer value in dB of the received signal strength. Every time a segment of the signal is found to have a positive slope, the histogram bins crossed by the segment are incremented by one. Thus, for example, if a given segment starts at -110 dB and rises to -105 dB in the next sample, the values of the bins representing -110, -109, -108, -107 and -106 dB are incremented by one. After the complete histogram has been assembled, the next step is preferably to determine the location on the dB scale of the noise floor level. Typically, the histogram will be bimodal in the presence of noise and a plurality of signals, with the first true local maximum indicating the noise floor. To find the noise floor, preferably the first local maximum starting from the minimum quantised dB value is selected. Because the histogram is created from sampled data affected by a random source it is possible to have numerous local maxima that do not reflect the true noise floor, as a result it is preferred to check if the local maximum is in fact the noise floor. In the most preferred embodiment, the method for the maximum value found includes the step of checking if a histogram value within Y dB is higher than the current maximum, where Y is a preferred number derived empirically from the characteristics of the system, most preferably between 2 and 3 dB. If the no value in the next Y db is greater than the value at the maximum then the current maximum abscissa is indicative of the noise floor level estimate per bins, however if the value is greater Y dB away from the maximum than it is at the maximum then it is clear that the noise floor has not been found. In the event that the maximum is not the noise floor, the next local maximum is preferably used and the Y dB test applied again until the test fails to produce a value higher than the value at the maximum. The obtained abscissa value is finally preferably corrected by a fixed offset of X dB, where X is preferably dependent on the selected window, to give the noise floor level estimate per bin for the processed block of data.

In a second preferred embodiment of the steps of creating and analyzing the histogram to estimate the noise floor, herein referred to as the penalized quasi log-likelihood method, a histogram of quantized frequency domain amplitudes is constructed with each bin 1 dB in size. The histogram is then transformed into a sorted vector of linear amplitude elements in

decreasing order. The sorted linear vector is then preferably reduced from the size MK to M by summing a group of K consecutive samples of the sorted vector. In this method M is typically much larger than K. A quasi log-likelihood function $L(k)$ is then preferably applied on the sorted group of M elements. The negative of the quasi log-likelihood function is

5 preferably added to a penalty polynomial function $p(k)$ giving an expression of the form
 $-L(k) + C p(k)$, where C is a constant, preferably empirically determined and most preferably near -2.6. The minimum of the $-L(k) + C p(k)$ function is preferably selected where the index k of the minimum amplitude is denoted by q_{NF} . If a line is drawn on the histogram at the index q_{NF} , it can be seen that the received signal can be divided into 2 groups, the first being the message subspace, where the information contained in the non-noise signal exists, and the noise subspace, where the random noise values exist. Finally, the noise floor level estimate per frequency bin is preferably obtained by dividing by K the noise floor level estimate per frequency bin is the average value of the power along the dimensions of noise subspace.

10
15 Frequency
A The application where the proposed methods have been used is in spectrum monitoring of communication radio signals. The methods are particularly suitable where fast scanning operations are necessary and an adaptive threshold value is required. They can also be used to automate lengthy manual noise level measurement procedures, and the resulting noise floor level estimates can be used as an indication of the spectrum quality.

20 As stated earlier, it is an object of the invention to derive an estimate in real-time of the noise floor level in a wideband signal preferably without knowledge of the signal content in terms of carrier frequency, bandwidth, and power, and using only a short-time signal. In the preferred apparatus in accordance with the invention, the wideband signal, containing a plurality of narrowband signals, is preferably received from the analogue output of a radio receiver and sent to a digitizer where a block of output data is generated over a time period. The sampled amplitudes, though discrete, are real valued and thus the signal is still considered continuous in amplitude. The block of output data is then windowed in a windowing module and sent to a channelizer where the wideband signal is represented by a collection of complex amplitudes at several discrete frequencies. The amplitudes of the

signals at the discrete frequencies are converted in a converter module to the log domain representation (decibels (dB) or decibel milliwatts (dBm)) and rounded to the nearest integer values. Thus the received signal is now represented by both discrete amplitude and frequency values.

5 Brief Description of the Drawings

Exemplary embodiments of the invention will now be further described with references to the drawings in which same reference numerals designate similar parts throughout the figures thereof, and wherein:

Figure 1 is a block diagram of the signal processing elements for the discrete level crossing rate (DLCR) embodiment of the method according to the invention;

Figure 2 is an example of the discrete level crossing rate with a noise only input signal;

Figure 3 is an example of the discrete level crossing rate with both noise and non-noise signals at the receiver's input

Figure 4 is an example of the behaviour of the DLCR method as a function of the occupancy;

Figure 5 is an example of experimental results showing the time behaviour with low and high signal frequency band occupancy for the DLCR method in the 800 MHz band;

Figure 6 is a block diagram of the signal processing elements for the penalized quasi log-likelihood (PLLM) embodiment of the method according to the invention;

Figure 7 is a plot of the quasi log-likelihood function and the associated penalizing polynomial;

Figure 8 is an example of the PLLM(k) function for a noise only signal, and a non-noise with noise signal;

5 Figure 9 is an example of the behaviour of the PLLM method as a function of the frequency band occupancy;

Figure 10 is an example of experimental results showing the time behaviour with low and high signal occupancy for the PLLM method in the 800 MHz band; and

Table 1 shows the offset for the DLCR method.

10 Detailed Description of the Preferred Embodiment

In the context of spectrum monitoring, the input signal contains several channels that are observed by a wide-band channelisation receiver. The receiver requires knowledge of the noise floor, in order to set a threshold that will produce a constant detection false alarm rate, when the only assumption is an input signal with additive noise normally distributed. For simplicity and speed, the receiver must estimate the noise floor level from the current 15 sampled information vector, which is provided by an analogue-to-digital (A/D) converter.

The time over which the input vector is sampled can vary depending upon the resolution required, the signal processing stages that follow and the physical constraints of moving the data to the subsequent stage. To ensure a fast scanning operation of the 20 channelised receiver, it is important not to use more samples to extract the necessary information from the input vector for the noise floor level estimation process than the channelised receiver needs. The presence and parameters of these non-noise signals may not be known by the monitoring equipment. The *a priori* knowledge of the receiver is minimal and is limited to the information that signals are not wider than the sampling bandwidth.

For a first embodiment of the present method, herein referred to as the Discrete Level Crossing Rate method, the signal path of the estimation method is presented in Figure 1. The signal to be analyzed 101 has previously been filtered by methods known in the art to prevent aliasing of the frequency bands after the sampling process. The input analogue signal 101 contains additive noise emanating from various sources and non-noise signals. A plurality of 5 analogue-to-digital converters 102 digitizes the signal and stores a series of samples in a buffer represented by the vector r 103. The vector r 103 is then windowed or polyphase filtered to produce another vector to be sent to a digital filter bank 104. The window choice may depend on specific requirements relative to a particular signal environment or on filter 10 bank design criteria. It is also important that any gain in the signal path (noise and non-noise signals) up to the filter bank output be identified, cancelled or compensated, so that the output of the filter bank is representing the true signal power at the antenna input. It is therefore necessary to ensure that the equipment is properly calibrated to avoid any biased estimation 15 of the noise floor level. In the digital domain, the filter bank is often implemented as a Fast Fourier Transform (FFT) to minimize computational complexity, but other transforms are also possible. The FFT is the preferred filter bank in this invention. The filter bank output represented by the vector z 105 is then transformed in a transforming module 106 to the integer log magnitude domain (dB or dBm) with a rounding operation on 20 times the logarithm base 10 of the magnitude of each complex point of the vector z . The implementation of the transformation step 106 is simplified by the rounding operation such 20 that the precision required in the log computation can be small. The resulting vector x 107 is then processed by the Discrete Level Crossing Rate (DLCR) module 108. The analogue Level Crossing Rate (LCR) function is well known in the signal processing community where it is mainly used in wave propagation and communication channel modelling to obtain 25 statistical information on the rate of change of fading processes on transmitted signals, and can be found in numerous standard texts including "Mobile Communications Engineering" W.C.Y. Lee, McGraw-Hill, 1982. The LCR is a second-order statistical operation that is time-dependent. It is usually sufficient to consider only the upward or the downward crossing in the evaluation of the function. The LCR is defined as the total number of expected 30 crossings of the signal envelope, at a given amplitude A , per second. The standard LCR operates on time domain signals and is a continuous level function with a continuous signal at the input. The current signal has a discrete-frequency discrete-amplitude representation. To

overcome this difficulty, the solution is to use a discrete amplitude version of the LCR in the form of a histogram containing the number of times a discrete decibel amplitude value is crossed in the upward direction. Therefore, with discrete-frequency signals, an approximation to the true LCR is built as a histogram of the number of occurrences of the positive crossing (positive slope) of an integer amplitude value in dB. For an additive Gaussian noise only signal, this procedure produces a function having a shape of the form .

$$n[Z = A] = \frac{\sqrt{2} \rho 10^{\frac{A}{10}}}{40\sqrt{\pi} \sigma^2 \log_{10}(e)} \exp\left(-\frac{10^{\frac{A}{10}}}{\sigma^2}\right),$$

where A is the integer level in dB at which the number of crossing is required, ρ is a correlation parameter of the frequency levels, and σ^2 is the variance of the additive complex Gaussian noise process in an FFT bin. The function above has a single maximum that can easily be verified to be located at $10 \log_{10}(\sigma^2)$. The parameter to estimate is the value of σ^2 and it can be estimated with the position of the maximum of the function $n[A]$ when only noise is present.

The situation where one or several non-noise signals are present with the additive noise needs additional signal processing. For this case, it is possible that the maximum of the function is not the noise floor level since the signal might cross a given amplitude more often if the observation bandwidth is highly occupied with non-noise signals. To avoid potentially large bias errors, the first local maximum, starting from the lowest dB values, is selected. To check if this local maximum is not due to a random fluctuation in the construction of a given discrete LCR histogram, it is then verified that there is no greater histogram value on the right side of the local maximum closer to Y dB relative to the current local maximum. This verification also stabilizes the method by preventing large overestimates of the noise floor level when the channel occupancy is high. The search for the appropriate maximum performed in search module 110, on the output of the Discrete Level Crossing Rate module 108, which is referred to as 109, can have different forms. The value of Y depends on the dynamic range of the non-noise signals. For this application, a value of 3 dB is used. Other

mechanisms to stabilize the method in presence of high occupancy narrowband signals are possible. The abscise value of the maximum, found in search module 110 is the noise floor level estimate 111 per FFT bin within an offset. The offset to be added to the abscise value depends on the window or the polyphase filter used. Table 1 gives the offset for four popular windows. Thus the noise floor level estimate $\hat{\sigma}^2$ per FFT bin in dB is

5 $\hat{\sigma}^2$ [dB] = position of appropriate maximum [dB] + Offset [dB].

The noise floor level estimates $\hat{\sigma}^2$ can be average for the same frequency band to reduce the variance if the noise floor level is known to vary very slowly.

To evaluate the performance of this method, simulation software was employed and a real-time implementation was developed to gather statistics about the noise floor level estimates. The performance in terms of the bias and the variance of the noise floor level estimate of the invented method was presented as a function of the channel occupancy, and as a function of the information signal power for a scenario of a specific non-noise signal. A sampled FM signal in the 139 to 144 MHz band was used as the transmitted signal in the simulation software. The channel bandwidth was set to 30 kHz and the windowed FFT filter bank used a Blackman window. The transmitted signals all have the same average power at the receiver input.

Figure 2 demonstrates the construction of the DLCR histogram for a particular vector \mathbf{x} when only a noise signal is present in the simulation data at the input. In Figure 2a, a Blackman window is used with 1024 filter bank outputs rounded to their nearest integer values. The resulting DLCR is presented in Figure 2b where the maximum 201 of the histogram is at -9 dB. Adding the 3 dB offset gives a noise floor level per FFT bin of -6 dB, the true value. Figure 3 demonstrates the construction of the DLCR histogram for a particular vector \mathbf{x} when a composite signal having a channel occupancy of 12.5% with $K = 4$ bins per channel is present with the noise signal in the simulation data at the input of the A/D. The signals have a signal power 20 dB above the noise floor level. In Figure 3a, a Blackman window is used with 1024 filter bank outputs rounded to their nearest integer values. The

resulting DLCR is presented in Fig. 3b where the maximum of the histogram is now at -8 dB
301. With the 3 dB offset, the noise floor level per FFT bin is of -5 dB. A bias of 1 dB is
introduced in this example. In general, the algorithm will introduce a small bias as a function
of the narrowband non-noise signal spectrum shape, the signal amplitude and more
5 importantly the channel occupancy. In the following, the number of channels M has been
fixed to 64 and the number of bins per channel K to 8. These numbers are typical of what is
expected in practice. Other choices of number of channels and number of bins per channel are
possible with different performance results. Figure 4 illustrates the performances for the
values of signal-to-noise ratios of 10, 20, and 30 dB in 30 kHz with the true value of the noise
10 floor level being at 0 dB. Each column bar is the average of 1000 independent estimates with
an error bar for the standard deviation. Figure 4 shows that the estimate bias is increasing
with the channel occupancy. It is within 3 dB of the true value up to 68.75 % of channel
occupancy 401. Figure 5 presents the real-time implementation results where 1000 estimates
were computed as a function of time for real-life signals. The signal of interest was over a 4.5
15 MHz bandwidth sampled at 5.12 MHz with 5 kHz bin resolution. The size of the vector r is
thus 1024. Two curves are presented: one where the centre frequency was at 848 MHz with a
low occupancy of approximately 3 %, and one where the centre frequency is 880.05 MHz
where the occupancy is about 54 % with an average signal power of 30 dB above the noise
floor level. Since the two bands are very close to each other, it is assumed that the noise floor
20 levels of both are the same. The algorithm shows a bias of about 2.2 dB for an occupancy
level of 54 % relative to the low channel occupancy band. The variance also increases when
the occupancy increases. In this particular case, the variance is tripled. The performance of
the method of the invention is thus very good considering that no *a priori* knowledge of the
signals is available to the receiver.

25 Another embodiment of the noise floor level estimation method of the invention is the
method presented in Figure 6, which is referred to herein as the penalized log-likelihood
method. The first seven steps of the method (601, 602, 603, 604, 605, 606, 607) are the same
as the first seven steps (101, 102, 103, 104, 105, 106, 107) of the DLCR method of the
embodiment presented in Figure 1. Thus, the received signal 601 to be analyzed is assumed
30 to have been filtered to prevent aliasing of the frequency bands after the sampling process. A
plurality of analogue-to-digital converters 602 digitizes the signal and stores a series of

samples in a buffer represented by the vector r 603. The vector r is then windowed or
5 polyphase filtered before being applied to a filter bank 604. It is also important that any gain
in the signal path (noise and non-noise signals) up to the filter bank output be cancelled or
compensated, so that the output of the filter bank is representing the true signal power at the
antenna input. It is therefore necessary to ensure that the equipment is properly calibrated to
10 avoid any biased estimation of the noise floor level. In the digital domain, the filter bank is
often implemented as a Fast Fourier Transform (FFT) because of its computational
efficiency. The filter bank output represented by the vector z 605 is then converted to the
integer log magnitude domain 606 with a rounding operation on the product of 20 times the
logarithm base 10 of the magnitude of each complex point of the vector z 605. Here again, a
15 requirement is that the noise power level be determined from a single FFT trace. The rounded
dB values in the vector x 607 are then used to build a standard histogram 608 with histogram
bins corresponding to the integer dB values. The histogram computed is then used to build an
ordered sequence 608 of amplitude dB values in decreasing order w 609. The vector w is then
20 of the same length and contains the same elements as the vector x but with elements in
decreasing order. Thus, the block of histogram and the decreasing ordered sequence 608 is
equivalent to a decreasing sort operation on the rounded dB values. The present approach is
much less complex in terms of number of operations required to perform the sorting than a
true sorting method. The decreasing sorted integer dB values are then transformed back to the
linear domain 610 with the transformation

$$v_i = 10^{w_i/10}$$

for each of the elements in the vector w 609. The decreasing sorted linear values are then
summed by groups of K 610 starting with the largest values. This grouping reduces the
dimension of the problem from MK to M . Another impact of the grouping operation is the
25 tendency of the group sample power to be de-correlated even with the use of a window 604,
because the decreasing ordered samples are not necessary from adjacent FFT bins. The
following block 612 is the computation of the penalized quasi-log-likelihood function with

the expression

$$\text{PLLM}(k) = -L(k) + Cp(k), k = 0, 1, \dots, M-1,$$

where $L(k)$ is a quasi log-likelihood function of the M sorted group elements, C is a constant and $p(k)$ is a penalty function. In this embodiment, the quasi-log-likelihood function is given

5 by

$$L(k) = K \ln \left[\frac{\prod_{i=k+1}^M l_i}{\left(\frac{1}{M-k} \sum_{i=k+1}^M l_i \right)^{M-k}} \right], k = 0, 1, \dots, M-1,$$

where l_i are the M decreasing values of the power in the grouped K bins 611, though it should be noted that a variety of log-likelihood, and quasi-log-likelihood functions could provide acceptable results. The function $L(k)$ would be a true log-likelihood function if l_i were the variance of independent and zero mean Gaussian random variables, but the elements of a received signal are correlated and thus are not independent nor are they necessarily zero mean Gaussian random variables. The penalty function $p(k)$ is added to the negative of the quasi log-likelihood function to form the function $\text{PLLM}(k)$ over which the minimum will be searched 614. The penalty function $p(k)$ is often related to an Information Theoretic Criteria 10 as described in "Detection of Signals by Information Theoretic Criteria" M. Wax and T. Kailath, IEEE Transactions Acoustic, Speech, Signal Processing volume 33, number 2, pages 15 387-392, April 1985. The penalty function in this embodiment, however, is different from those derived earlier in that the coefficients l_i are not truly statistically independent, and that the intent is not to determine the full model order as in the reference cited above. For this 20 embodiment, the penalty function is a second order polynomial given by

$$p(k) = C \left[3.76 \left(\frac{M-1-k}{M-1} \right)^2 + 1.43 \left(\frac{M-1-k}{M-1} \right) \right] MK.$$

Other polynomials or functions are possible, but this second order polynomial is adequate for a good performance and complexity trade off. The polynomial function is scaled by a constant C to avoid erroneous detection due to the variance in the data. The value of C has been determined empirically and may vary depending on the final goal of the design or the signal environment at the input. For this application, $C = -2.6$ is the preferred value. The only assumption to determine the above polynomial is that the FFT coefficients are distributed according to a Gaussian probability density function. If the additive noise time data samples do not have a Gaussian distribution the linear combination of the FFT will tend to transform them into variables having a probability distribution close to Gaussian variables for practical and large size of MK . The method is therefore robust to non-ideal conditions where the additive noise might not be truly Gaussian. The scale factor C gives a degree of freedom in the optimization process. The output of function $\text{PLLM}(k)$ is then searched in step 614 for the index of its minimum q_{NF} defined as

$$15 \quad q_{\text{NF}} = \arg \min_k \{-L(k) + Cp(k)\}.$$

The above expression means that q_{NF} is the index at the minimum of the function $\text{PLLM}(k)$. The noise floor estimate per K FFT bins is the average of the $M - q_{\text{NF}} - 1$ smallest numbers of the sorted vector values l_i . The noise floor estimate 615 per FFT bin is then given by

$$20 \quad \hat{\sigma}^2 = \frac{1}{K(M - q_{\text{NF}} - 1)} \sum_{i=q_{\text{NF}}+1}^{M-1} l_i.$$

The noise floor level estimates $\hat{\sigma}^2$ can be averaged for the same frequency band to reduce the variance if the noise floor level is known to vary very slowly.

To evaluate the performance of the PLLM algorithm, simulation software was employed and a real-time implementation was developed to gather statistics about the noise floor level estimates. The performance in terms of the bias and the variance of the noise floor level estimate of the invented method is presented as a function of the channel occupancy and as a function of the information signal power for a scenario of a specific signal. A sampled FM signal in the 139 to 144 MHz band is used as the transmitted signal in the simulation software. The channel bandwidth is set to 30 kHz and the windowed FFT filter bank used a Blackman window. The transmitted signals all have the same average power at the receiver input.

Figure 7 illustrates the normalized quasi log-likelihood function 701 and the second-order polynomial approximation 702 with the scale factor $C = 1$ for a Gaussian noise only input signal. As the figure shows, the two curves are similar over the range of the index k . Figure 8 presents the $\text{PLLM}(k)$ function for a noise only, and a non-noise plus noise signal. It can be seen that the index of the minimum of the $\text{PLLM}(k)$ function is indicative of the signal subspace dimension. Since the remaining subspace includes noise only components, then the average is an estimate of the noise floor level. In the following, to reduce the amount of simulation results to present, the number of channels M has been fixed to 64 and the number of bins per channel K to 8. These numbers are typical of what is expected in practice. Other choices of number of channels and number of bins per channel are possible with different performance results. Figure 9 illustrates the performances for the values of signal-to-noise ratios of 10, 20, and 30 dB in 30 kHz with the true value of the noise floor level being at 0 dB. Each column bar is the average of 1000 independent estimates with an error bar representing the standard deviation. Figure 9 shows that the estimate bias increases with the channel occupancy. The estimate is within 2 dB of the true value up to 84.375 % of channel occupancy 901. Figure 10 presents real-time implementation results where 1000 estimates of the noise floor level were computed using real-life signals. The signal of interest was over a 4.5 MHz bandwidth sampled at 5.12 MHz with 5 kHz bin resolution. Two curves are presented: one where the centre frequency was at 848 MHz with a low occupancy of approximately 3 %, and one where the centre frequency is 880.05 MHz where the occupancy is about 54 % with an average signal power of 30 dB above the noise floor level. Since the two bands are very close to each other, it is assumed that the noise floor levels of both are the

same. The algorithm shows a bias of about 1.15 dB for an occupancy level of 54 % relative to the low channel occupancy band. The variance also increases when the occupancy increases. In this particular case, the variance is almost tripled. The performance method is thus very good considering that no *a priori* knowledge of the signals is available.

5 The exemplary embodiments described above are provided for the sake of clarification and are not intended to limit the scope of the invention. Of course, numerous variations and adaptations may be made to the above-described embodiments of the invention, without departing from the spirit of the invention, which is defined in the claims.

10 Changes and modifications in the specifically described embodiments can be carried out without departing from the scope of the invention which is intended to be limited only by the scope of the appended claims.

Window	Offset (dB)
Rectangular	2
Hanning	2.75
Blackman	3
Flat-top	3.2

Table 1